

# Outsmarting Network Security with SDN Teleportation

[Preliminary Report]

Kashyap Thimmaraju<sup>1</sup> Liron Schiff<sup>2</sup>

Stefan Schmid<sup>3</sup>

<sup>1</sup> TU Berlin, Germany

<sup>2</sup> Tel Aviv University, Israel

<sup>3</sup> Aalborg University, Denmark

**Abstract**—Software-defined networking is considered a promising new paradigm, enabling more reliable and formally verifiable communication networks. However, this paper shows that the separation of the control plane from the data plane, which lies at the heart of Software-Defined Networks (SDNs), introduces a new vulnerability which we call *teleportation*. An attacker (e.g., a malicious switch in the data plane or a host connected to the network) can use teleportation to transmit information via the control plane and bypass critical network functions in the data plane (e.g., a firewall), and to violate security policies as well as logical and even physical separations. This paper characterizes the design space for teleportation attacks theoretically, and then identifies four different teleportation techniques. We demonstrate and discuss how these techniques can be exploited for different attacks (e.g., exfiltrating confidential data at high rates), and also initiate the discussion of possible countermeasures. Generally, and given today’s trend toward more intent-based networking, we believe that our findings are relevant beyond the use cases considered in this paper.

## I. INTRODUCTION

Computer networks such as datacenter networks or the Internet have become a critical infrastructure [14]. Not only a large fraction of the economic activity critically depends on the availability of such networks, but also governments increasingly rely on existing and shared infrastructures, mainly for their cost benefits [27].

This dependency on public and shared infrastructures raises concerns. While the Internet has certainly been a huge success, and over the last years, there has been much innovation on the higher network layers (e.g., application layer) and the lower network layers (e.g., data-link and physical layer), the Internet core suffers from ossification [6]. In particular, it is questionable whether today’s network technology is sufficient to ensure essential security, resilience and dependability properties. For instance, today’s Internet does not provide any means of path control, and we are still struggling to render routing protocols more secure [30].

Software-Defined Networking is a novel networking paradigm which promises to enable these necessary innovations, also in terms of security, through its decoupling and consolidation of the control plane, its formally verifiable policies [20], [21], [31], [34], as well as by introducing new functionality [29], [33], [36], [37].

However, Software-Defined Networks (SDNs) also introduce new security challenges. In particular, we in this paper study




Fig. 1: Illustration of teleportation: malicious switches (with *red horns*) exploit the control platform for hidden communication, possibly bypassing data plane security mechanisms such as a firewall.

threats introduced by an *unreliable south-bound interface*, i.e., we consider a threat model in which switches or routers do not behave as expected, but rather are malicious, and e.g., contain hardware backdoors [4]. While many existing network security and monitoring tools rely on the trustworthiness of switches and routers, this assumption has become questionable: attackers have repeatedly demonstrated their ability to compromise switches and routers [1], [5], [26], thousands of compromised access and core routers are being traded underground [25], networking vendors have left backdoors open [2], [3], national security agencies can bug network equipment [4], hacker tools to scan and eventually exploit routers with weak passwords, default settings are openly available on the Web, etc. However, the impact of malicious hardware is not well understood and underexplored today.

In particular, this paper shows an outsourced and consolidated control plane, as it lies at the heart of the SDN paradigm, introduces an opportunity for *teleportation*: malicious SDN switches may transmit information via the logically centralized software control plane, *completely* bypassing data plane elements (such as other switches, middleboxes, etc.). By violating logical and even physical separations, teleportation can constitute a serious security threat. For example, teleportation could be used by one malicious switch to discover

(and communicate information to) other malicious switches, bypassing security checks in the data plane. As we will show in this paper, teleportation can also be exploited by malicious hosts, triggering (benign) switches to teleport information for them.

We argue that teleportation can be seen as a flexible communication channel which constitutes a threat in various situations, for example (see Figure 1):

- 1) **Bypassing critical network components:** By implicitly communicating information via the control plane, it is possible to circumvent critical network components, such as switches, middleboxes or policy enforcers located in the data plane. For example, teleportation can in principle be used to bypass middleboxes performing security checks (e.g., network intrusion detection systems), middleboxes in charge of billing (e.g., a Radius server), or QoS enforcers (e.g., a leaky bucket policer).
- 2) **Rendezvous and attack coordination:** While already a single malicious switch, for example located *inside* an enterprise network, may cause significant harm and violate basic security policies, the situation becomes worse if multiple malicious switches cooperate [13]. Malicious or Trojan switches (e.g., switches containing a hardware/software backdoor) may use teleportation as a rendezvous protocol, to discover each other, and subsequently coordinate an attack.
- 3) **Exfiltration:** Teleportation can also be used to exfiltrate sensitive information between networks that have no data plane connectivity.
- 4) **Eavesdropping and data tampering:** Particularly serious threats are introduced if malicious hosts and switches collude. For instance, in a scenario with collusion, teleportation can be used for eavesdropping. We show that a malicious switch and host can carry out a man-in-the-middle attack that serves benign hosts with malicious web pages.

Teleportation can be difficult to detect: the teleported information follows the normal traffic pattern of control communication, not between switches directly but *indirectly* between any switch and the controller. Moreover, the teleportation channel is *inside* the (typically encrypted) OpenFlow channel. Accordingly, it cannot easily be detected with modern Network Intrusion Detection Systems (NIDSs), even if they operate in the control plane.

#### A. Our Contributions

This paper makes the following contributions:

- 1) We identify a new vulnerability which targets the very core of the software-defined networking paradigm, namely the separation of the control and the data plane. In particular, we consider the threats introduced by a malicious data plane. Indeed, recent incidents related to the trustworthiness of routers and switches, indicate that our threat model is a relevant one.
- 2) We model and characterize possible teleportation channels theoretically.

- 3) We recognize and demonstrate four teleportation techniques in software-defined networks utilizing state-of-the-art OpenFlow controllers, in particular ONOS [38] among others.
- 4) We present and demonstrate multiple simple and more sophisticated attacks. In particular, we show that teleportation can also be exploited by malicious hosts in scenarios where all switches are benign.
- 5) We initiate the discussion of possible countermeasures. In particular, we propose to combine intrusion detection with waypoint-enforcement, forcing *Packet-out* messages (from controller to switches) to pass through the waypoint if mandated by a security policy.

We have already notified the open source community about some of the issues reported in this paper, and first actions have been taken (see CVE-2015-7516).

More generally, in the light of today's trend toward more intent-based networking, we believe that our work touches a topic whose relevance may increase in the near future and go beyond the use cases considered in this paper.

#### B. Paper Organization

The remainder of this paper is organized as follows. Section II introduces the necessary background on OpenFlow and SDN. Section III introduces our threat model, and Section IV characterizes possible teleportation channels. Section V describes teleportation techniques; based on these techniques, we demonstrate and discuss different attacks in Section VI. Section VII initiates the discussion of countermeasures. After reviewing related work in Section VIII, we conclude our work in Section IX.

## II. PRELIMINARIES

This paper considers Software-Defined Networks (SDNs) which outsource and consolidate the control over the network switches to a logically centralized software controller. The separation of the control and data plane has the potential to simplify the network management, as many networking tasks are inherently non-local. Moreover, SDN and especially its de facto standard protocol, OpenFlow, introduce interesting new flexibilities, e.g., in terms of traffic steering: routes may not necessarily be destination based, and can depend on layer-2, layer-3 and even layer-4 properties of the packets.

At the heart of an SDN lies a control software, running on a set of servers. These controllers receive information and statistics from switches, and depending on this information as well as the policies they seek to implement, issue instructions to the switches.

OpenFlow follows a match-action paradigm: The controllers install rules on the switches which consist of a match and an action part; the packets (i.e., flows) matching a rule are subject to the corresponding action. That is, each switch stores a set of (flow) tables which are managed by the controllers, and each table consists of a set of flow entries which specify expressions that need to be matched against the packet headers, as well as actions that are applied to the packet when a given

expression is satisfied. Possible actions include dropping the packet, sending it to a given egress port, or modifying its header fields, e.g., adding a tag. The match-action paradigm is attractive as it simplifies formal reasoning and enables policy verification.

By default, if a packet arrives at a switch and does not match an existing rule, the packet (usually without payload if the switch supports packet buffering) is forwarded to the controller. This event is called a *Packet-in*. Upon a *Packet-in* event, the controller can decide how to react to packets of the corresponding type, and add/delete/modify flow rules accordingly issuing *Flow-mod* messages to the switch (and maybe to other switches proactively on this occasion as well). A controller can also decide to send out a packet explicitly from a switch, issuing a so-called *Packet-out* command to the switch.

An attractive alternative to the hop-by-hop installation of new flows, reacting to a new packet repeatedly along the path (multiple *Packet-ins*), is the so-called “pave-path technique”: once the controller receives a first *Packet-in* event from some switch, it proactively updates the other switches along the path. Such an “intent-based” controller behavior can render the reaction to network event and set up of host-to-host/network connectivity (according to current policies) more efficient.

While SDNs are logically centralized, the control plane can be physically distributed, e.g., for fault-tolerance or performance reasons. Accordingly, OpenFlow supports multiple controllers for a single switch. The controllers and switch exchange *Role-request* and *Role-reply* messages respectively to assert the various roles (*Master*, *Equal* and *Slave*). There may be only one *Master* controller for a given switch while multiple *Equal* and *Slave* controllers are permitted.

The OpenFlow standard specifies basic security mechanisms. For example, the communication between the controller and switch can be authenticated and encrypted, using TLS over TCP/UDP.

Finally, we note that although some of our techniques are generally applicable in networks separating the control plane and the data plane, while others exploit OpenFlow specific features, when clear from the context, in this paper we will treat SDN and OpenFlow as synonyms.

### III. THREAT MODEL

We in this paper consider a threat model where OpenFlow switches, hosts, or both, may not behave correctly but be malicious.

We do not place any restrictions on what a malicious switch can and cannot do. For example, a malicious switch can fabricate and transmit any type of OpenFlow message, it can arbitrarily deviate from the OpenFlow specification, and it can even use multiple identities, all at the risk of being detected. However, the malicious switch cannot choose where it will be placed in the network. In order to collude, the malicious switches have been programmed to recognize some data and/or timing pattern. Similarly, we do not place any restrictions on what a malicious host can and cannot do. For

example, a malicious host may masquerade its Media Access Control (MAC) and/or Internet Protocol (IP) addresses, use an incorrect gateway, falsify Address Resolution Protocol (ARP) requests/responses, and so on. The attacker could also be an insider, i.e., an authorized user who intends to subvert his/her current organization. We also consider the case where malicious hosts and switches collude. We assume that an attacker has sufficient resources and know-how to compromise hosts/switches and therefore do not concern ourselves with how the host/switch is compromised.

The OpenFlow controller and its applications on the other hand are trusted entities and are available to the switches: for example, they are based on static and dynamic program analyses. The OpenFlow channel is reliable and may be encrypted.

### IV. MODELING TELEPORTATION

With these concepts in mind, we now model and characterize a novel threat called *teleportation* which targets the heart of SDNs: the outsourcing and consolidation of control over multiple data plane elements. In particular, we argue that we can see an OpenFlow controller as a “reactor”: it reacts (in a best-effort and timely manner) to events generated by the network operator, the OpenFlow switches, and timeouts; as a response, the controller sends OpenFlow commands to switches. Accordingly, we argue that the following 3-stage functionality is fundamental in the SDN paradigm.

- 1) **Switch to controller:** A source switch intentionally or unintentionally sends modulated information to the controller (e.g., by adding specific events, delaying existing events, etc.).
- 2) **Controller to switches:** The controller reacts to the received events, by sending commands to one or multiple other switches.
- 3) **Destination processing:** A destination switch processes incoming commands from the controller. In case of a malicious switch, the switch may search for some message properties, temporal patterns, etc., and hence infer the information modulated by the source, or by simply forward the information (to a potentially malicious host).

Based on this controller model, we can identify two kinds of teleportation channels:

- **Explicit teleportation:** The teleported information actually appears in the messages exchanged. The message may for example contain steganographic contents.
- **Implicit teleportation:** The teleportation relies on modulating information implicitly. For example, it is based on timing (e.g., message transmissions are delayed according to some pattern) or it is based on shared resources, whose availability is changed over time (e.g., leveraging mutual exclusion).

### V. TELEPORTATION TECHNIQUES

Having established a conceptual model of teleportation, we next present techniques that can realize teleportation in today’s

SDNs. In particular, we have identified the following three fundamental SDN functionalities which can be exploited for teleportation:

- 1) **Flow (re-)configurations:** In an SDN, a controller needs to react to various data plane events (such as so-called *Packet-ins* in OpenFlow or link failures), and configure and reconfigure flows and paths accordingly. Triggering and exploiting such events can be used for teleportation.
- 2) **Switch identification:** In an SDN, switches are responsible for introducing and uniquely identifying themselves to the controller. This is required as policies are often specific to the switch. Unique switch identifiers are also necessary to correctly construct and enforce policies on the switches and in the controller. We will show that such switch identification mechanisms can be exploited for teleportation.
- 3) **Out-of-band forwarding:** An SDN controller must not only be able to receive events and control packets from switches, but also to instruct switches to forward specific messages. This basic functionality in SDNs can be exploited by a malicious switch or host to forward entire packets via the controller.

In the remainder of this section, we will discuss these teleportation techniques in more detail in turn.

#### A. Flow (Re-)Configurations

We distinguish between two types of flow re-configuration functions: *path update* and *path reset*.

a) **Path Update:** The first technique based on flow reconfiguration we consider is *path update*. Path updates are a fundamental controller functionality, and come in the form of different controller features such as *Mobility*, *VM Migration* or simply *MAC Learning*. The basic scheme is as follows: A controller typically maintains some mapping of which hosts (MAC addresses) are connected to which ports (on the switch). If a host suddenly appears on another switch, the controller installs new flows for the host on the new switch, and also deletes the corresponding flow rules on the old switch. We define this type of installation and deletion of flows by the controller on switches as *path update*. Specifically, a path update involves the use of *Packet-in*, *Flow-mod* and *Packet-out* messages. Malicious switches can use path update for implicit teleportation.

For the teleportation with path update, a switch triggers the deletion of rules at other switches. Malicious switches can teleport information between themselves by prompting path updates for the same host using *Packet-ins*. Note that during a path update, the *Packet-out* is sent to the destination reported in the *Packet-in* which may generate data plane traffic. To prevent data plane traffic, the malicious switch can use a destination host that is connected to itself (so that the *Packet-out* is sent back to it). The message sequence pattern for path update teleportation is shown in Figure 2.

We can summarize the scheme presented so far with the following abstract steps: a switch  $s1$  announces  $X$ , a switch  $s2$  announces  $X$  thereby stealing  $X$  from  $s1$ , where stealing is

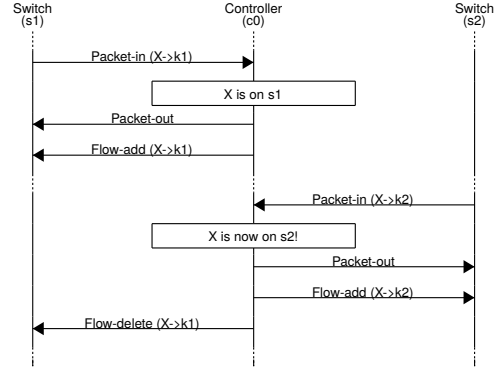


Fig. 2: Message sequence pattern for path update teleportation. Switch  $s2$  teleports information to  $s1$  when  $s1$  receives the Flow-delete ( $X \rightarrow k1$ ) message from controller  $c0$ .

**process incoming OpenFlow message :**

```

on start teleportation
  - announce all  $\{X_{i,j}\}_{j \in [m]}$  and  $\{X_{j,i}\}_{j \in [m]}$ 
on received Flow-delete for  $X_{i,j}$  for some  $j \in [m]$ 
  - announce  $X_{i,j}$ 
  - add  $j$  to Discovered_Switches

```

**Algorithm 1:** Generalized pseudo-code executed by switch  $s_i$  to teleport information using path update.

detected by the “victim” ( $s1$ ). Also note that announcing is possible once some host which is connected to the malicious switch (e.g.,  $k1$  at  $s1$ ) is learned by the controller.

Based on these basic steps, we can generalize our scheme for  $m$  malicious switches. Each malicious switch,  $s_i$ , with  $i \in [m]$ , should implement an event handler (see Algorithm 1), in addition to the normal (non-malicious) behavior. We assume that all switches are programmed with the same list of  $m^2$  special MAC addresses  $\{X_{i,j}\}_{i,j \in [m]}$  (the *pre-shared secrets*). Note that once switch  $i$  discovered switch  $j$ , it can contact it by sending packets with source  $X_{i,j}$  and destination address  $X_{j,i}$ .

b) **Path Reset:** We next discuss a second flow reconfiguration based teleportation which we refer to as *path reset*. Recall that at the heart of any SDN controller lies the functionality to set up host-to-host/network connectivity, according to the network policy (e.g., defining constraints such as bandwidth, link type and waypoints), which is translated into device level configurations (e.g., flow rules). The “pave-path technique” is an attractive alternative to the hop-by-hop installation of new flows: once the controller receives a first *Packet-in* event from some switch, it proactively updates the other switches along the path.

In order to provide high availability, a controller also monitors the network state and makes necessary changes, such as rerouting or resetting flows on switches, when needed (e.g., due to a link failure). For example, triggered by a *Packet-in* event, a controller may learn that (parts of) the path may no longer be available, and hence initiates the




Fig. 3: Message sequence pattern for path reset teleportation. Switch  $s_1$  teleports information to  $s_2$  and  $s_3$  via *Flow-add* messages sent by the controller  $c_0$ .

reconfiguration/repair of the path. We define the reinstallation of flows by the controller on switches along a path as *path reset*. Accordingly, the path reset technique involves *Packet-in*, *Flow-mod* and *Packet-out* messages.

Malicious switches may use path reset for implicit teleportation: If the controller resets the complete path between hosts when it receives a *Packet-in* from a switch that ignores the flow rule, then information can be communicated. By doing this at multiple and specific times, a malicious switch can teleport information to other malicious switches along the path. Figure 3 illustrates the message sequence pattern for teleportation using path reset.

### B. Switch Identification

This teleportation type exploits the fact that a switch typically must uniquely identify itself whenever it connects to the controller. For example, in OpenFlow this is usually done using the Datapath-ID (DPID) field in the *Features-reply* message. We define two switches attempting to use the same DPID to connect to the same logical controller as *switch identification*. The outcome can be used for implicit teleportation.

Three basic ways an OpenFlow controller can react to using the same DPID are as follows:

- 1) The controller denies the second switch a connection.
- 2) The controller terminates the first switch and connects to the second.
- 3) The controller accepts both switches but sends them different *Role-request* messages.

With any of the above outcomes a switch can infer the (mis)use of the same DPID by another switch. By using a-priori configured single or multiple DPID values, a pair of malicious switches can establish teleportation. For example, consider the message sequence pattern in Figure 4, and assume that first switch  $s_1$  tells controller  $c_0$  that its DPID is 1. At a later time, switch  $s_2$  tells  $c_0$  that its DPID is 1. At this point,  $c_0$  does not allow  $s_2$  to connect with DPID 1. Since  $c_0$  denied  $s_2$  to connect with DPID 1,  $s_1$  teleported information to  $s_2$  via  $c_0$ . With a similar message sequence pattern, the second outcome can be used for teleportation as well.

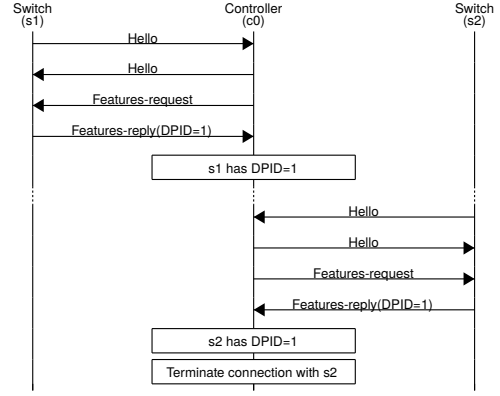


Fig. 4: Message sequence pattern for *switch identification* teleportation when the controller denies the second switch a connection. When  $s_2$ 's connection is terminated,  $s_1$  successfully teleports information to  $s_2$ .

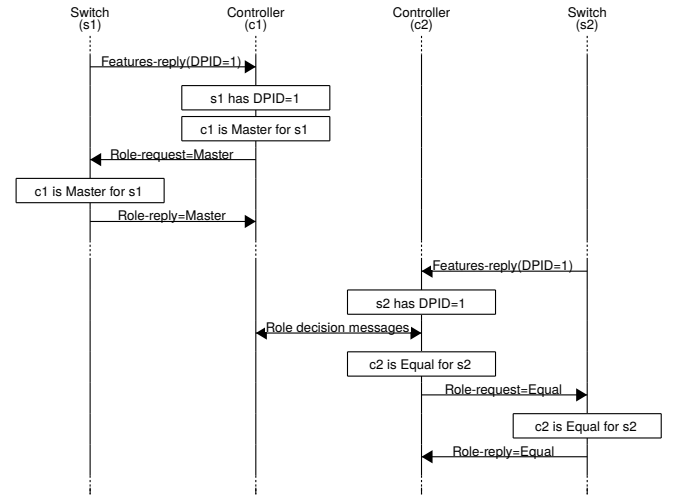


Fig. 5: Message sequence pattern for *switch identification* teleportation when controllers  $c_1$  and  $c_2$  send different *Role-request* messages to  $s_1$  and  $s_2$  respectively. When  $s_1$  receives the *Role-request=Master* message whereas  $s_2$  receives the *Role-request=Equal* message. In this manner  $s_1$  teleports information to  $s_2$  when  $s_2$  received the *Role-request=Equal* message.

Interestingly, switch identification is not limited to scenarios with a single controller: we have found additional threats in the presence of *distributed control planes*. Moreover, we can generalize the first *switch identification* outcome to scenarios with  $m$  malicious switches, see the event-handler algorithm, Algorithm 2. The other two outcomes discussed can also be seen as event-handler algorithms.

### C. Out-of-band Forwarding

The third and potentially most powerful teleportation technique is called *out-of-band forwarding*. It is an example of explicit teleportation. Out-of-band forwarding exploits the fact that an SDN controller is typically connected to multiple

**process** connect to OpenFlow controller :

```

on Features-request message from controller
├ - announce DPID  $\{X_{i,j}\}_{j \in [m]}$  in Features-reply
  message
on Controller denies connection to announced DPID  $X_{i,j}$ 
  for some  $j \in [m]$ 
├ - add  $j$  to Discovered_Switches

```

**Algorithm 2:** Generalized pseudo-code executed by switch  $s_i$  for *switch identification* teleportation when the controller denies the second switch a connection.

Technique	Type	Threat
Flow (Re-)Configuration	Implicit	covert communication and coordination
Switch Identification	Implicit	attack coordination
Out-of-band Forwarding	Explicit	exfiltration, firewall/NIDS bypass and man-in-the-middle

TABLE I: Summary of teleportation techniques, types and associated attacks.

switches: accordingly, a packet from one switch can potentially reach multiple other switches in the network via the control plane. Out-of-band forwarding involves *Packet-in* and *Packet-out* messages with the possible side effect of *Flow-mod* messages. Out-of-band forwarding could for example include the complete Ethernet frame (typically 1500 bytes), and can even serve as a “multicast service”. Out-of-band forwarding can be a serious threat to network security, not only because malicious traffic can bypass critical security functions in the data plane, but also because it can be exploited both by switches and hosts.

A summary of our teleportation techniques is shown in Table I along with the type of teleportation and the associated OpenFlow messages.

## VI. SWITCH- AND HOST-BASED ATTACKS

Let us now demonstrate how the identified teleportation techniques could be exploited to carry out specific attacks. In particular, we show how teleportation may be exploited:

- 1) to bypass security critical network functions such as firewalls and Network Intrusion Detection Systems (NIDS);
- 2) as a rendezvous protocol for malicious switches;
- 3) to exfiltrate sensitive data from remote locations;
- 4) to conduct a man-in-the-middle (mitm) attack.

Along the way, we also present a novel denial-of-service (dos) attack (published as CVE-2015-7516).

Before presenting the attacks in more detail, we report on the setting in which we tested the attacks.

### A. Setup

We verified all our attacks in a virtual machine, using Mininet-2.2.0, Open vSwitch-2.0.1 in the data plane. For

the control plane we used ONOS-1.3.1 as it was the state-of-the-art. At the time of our experimentation Floodlight, OpenDaylight Lithium-SR2 and RYU v3.27 still did not support the intent based framework. Indeed our experiments showed that they were only vulnerable to a subset of the attacks (e.g., switch identification, out-of-band-forwarding<sup>1</sup>) For packet generation we use ping and nmap-6.40. We use *ebtables* v2.0.10-4 (December 2011) as our transparent firewall and *Snort* version 2.9.6.0 GRE (Build 47) as our NIDS. We modified code developed by *austinmarton* [8] to set the Ethertype field in an Ethernet frame. *ettercap* 0.8.0 is used with a custom HTTP filter for the mitm attack.

### B. Bypassing Critical Network Functions

We believe that the possibility to bypass network elements is a serious threat in modern computer networks. For example, many network policies today are defined in terms of adjacency matrices or big switch abstractions, specifying which traffic is allowed between an ingress port  $s$  and an egress network port  $t$  [19]. In order to enforce such a policy, traffic from  $s$  to  $t$  needs to traverse a middlebox instance (waypoint) inspecting and classifying the flows. The location of every middlebox may be optimized, but is subject to the constraint that the route from  $s$  to  $t$  should always go via the waypoint.

*Firewall and NIDS:* In order to demonstrate how a firewall may be circumvented by hosts (or switches), we set up Mininet and ONOS as shown in Figure 6. The switches do not have flow rules for  $k1$  and  $k2$  to communicate. The firewall *fw1* prevents hosts on the left to communicate with hosts on the right and vice-versa. ONOS has the *Intent Reactive Forwarding (ifwd)* application enabled. *ifwd* uses the reactive “pave-path technique” (discussed above) to install flows in the switches. By default, the *ifwd* application establishes host-to-host connectivity when it receives a *Packet-in* for which no flows exist.

We send a ping packet from  $k1$  to  $k2$ . Despite the presence of the firewall,  $k1$  receives the reply from  $k2$  using out-of-band forwarding teleportation. In the absence of out-of-band forwarding teleportation, the packet would have been dropped by *fw1*.

Indeed, in this case out-of-band forwarding teleportation has the side effect of installing flows on  $s1$  and  $s2$  for  $k1$  and  $k2$  to communicate, preventing further out-of-band forwarding teleportation. By masquerading its MAC address,  $k1$  can teleport more data to  $k2$  via out-of-band forwarding teleportation.

Similar to the firewall scenario, we can also use out-of-band forwarding teleportation in the presence of *Snort*, an NIDS. In particular, we can generate attack traffic using nmap to conduct TCP flag attacks or even port scans. Indeed, by masquerading the source MAC address, one can effectively carry out a wide enough port scan without having the scan pass through the firewall and being detected by the latter.

By replacing the firewall we previously described with *Snort*, we use nmap from  $k1$  to carry out a TCP port scan on  $k3$  using out-of-band forwarding teleportation. By inspecting the

<sup>1</sup> <https://goo.gl/FN9ULQ>




Fig. 6: An SDN topology with OpenFlow switches  $s1$  and  $s2$  and an OpenFlow controller  $c0$ .  $k1$  and  $k3$  are connected to  $s1$  while  $k2$  and  $k4$  are connected to  $s2$ .  $s1$  and  $s2$  are separated by a firewall  $fw1$  that denies hosts on  $s1$  to communicate with hosts on  $s2$  and vice-versa.  $k1$  can use *out-of-band forwarding teleportation* to transfer data to  $k2$ , bypassing  $fw1$ .




Fig. 7: An SDN topology of OpenFlow switches  $s1$ ,  $s2$ ,  $s3$  and  $s4$ , OpenFlow controller  $c0$ . Hosts  $k1$  and  $k3$  are connected to  $s1$  and  $k2$  and  $k4$  are connected to  $s2$ .  $c0$  has installed flows on  $s1$ ,  $s3$  and  $s2$  so that  $k1$  and  $k2$  can communicate bi-directionally. Teleportation traffic is via  $c0$ .

alerts in *Snort* we verified that no alerts were generated for the port scan.

Note that the host-to-host connectivity setup involves a *Packet-in* and *Flow-mod* messages whereas the out-of-band forwarding teleportation only involves *Packet-in* and *Packet-out* messages with the side effect of *Flow-mod* messages. Therefore, security policy enforcers that do not inspect and correlate *Packet-in* with *Packet-outs*, will miss out-of-band forwarding teleportation based attacks. Of course, violating *Flow-mods* may eventually be detected, but only after the data has been teleported.

### C. Rendezvous and Malicious Switch Discovery Protocol

We next consider a rendezvous protocol in which malicious switches wish to discover one another. A rendezvous or discovery protocol can be also seen as a precursor to a much more damaging attack such as a denial-of-service, man-in-the-middle (mitm) or exfiltration.

Teleportation can be an attractive solution: a rendezvous protocol can rely on steganography i.e., embedding patterns in teleported benign information or modulating patterns in legitimate messages. Without teleportation, by going through the data plane directly, the malicious switches risk detection.

We show how three of our techniques, namely path update, path reset and switch identification teleportation may be used as a rendezvous protocol for malicious switches.

1) *Path Update*: To demonstrate a rendezvous with path update teleportation, we set up Mininet and ONOS as shown in Figure 7. Instead of instrumenting code for the malicious switches, we keep them as simple Open vSwitches and we defined dedicated Mininet hosts ( $k3$  and  $k4$ ) for each of them. We use the dedicated hosts ( $k3$  and  $k4$ ) to generate the packet that the malicious switch sends as a *Packet-in* to the controller. The *host mobility* and *ifwd* applications are enabled on ONOS. The controller has already installed flows for  $k1$  to  $k2$  and vice-versa. Accordingly, we use  $k4$  connected to  $s2$ , to send  $k2$

a packet using  $k1$  as the source MAC address. This triggers the controller to issue *Flow-mod* commands to  $s1$ ,  $s2$  and  $s3$ .  $s2$  thereby teleported its presence to  $s1$ .

By inspecting the flows on the switches, we verified the successful path update teleportation:  $s2$  was able to cause a flow deletion in  $s1$  without exchanging any packets with  $s1$  directly (except for a normal flow in the past).

Note that path update may trigger alerts in systems that keep track of moving MAC addresses by inspecting *Packet-in* and *Flow-mod* messages. In such cases, many moving MAC addresses may introduce suspicious activity within the network. Also worth noting is that port-based security (that associates MAC addresses with specific ports) is not applicable in the presence of malicious switches.

2) *Path Reset*: To demonstrate that path reset teleportation can be used as a rendezvous protocol, we consider the same setup as outlined in Section VI-C1. We modulate traffic between  $k1$  and  $k2$  using ping packets with 100 microseconds intervals. Instead of manipulating the Open vSwitch code for sending a *Packet-in* for an existing flow from  $s1$ , we simply remove the flow for  $k1$  to  $k2$  on  $s1$ , using the `ovs-ofctl del-flow` command. This causes  $s1$  to send  $c0$  a *Flow-removed* message which triggers the controller to add the flow back onto  $s1$ . But due to the high rate of ping traffic, at least one packet triggers a table-miss before  $s1$  adds the flow and a *Packet-in* is sent to  $c0$ .

When  $c0$  receives the *Packet-in* it sends the packet to  $s2$  directly as a *Packet-out*, bypassing  $s3$ , and then sends *Flow-mods* to  $s1$ ,  $s2$  and  $s3$  resetting the bi-directional path between  $k1$  and  $k2$ . By checking the lifetime of the flow rules on  $s1$ ,  $s2$  and  $s3$  we verified that path reset teleportation succeeded. In this manner,  $s1$  teleported its presence to  $s2$  by having the controller send *Flow-mod* commands for existing flow rules.

Note that such an attack works in the presence of topology spoofing defenses [11], [15] as the *Packet-in* and *Flow-mod*






Fig. 8: An SDN topology with independent OpenFlow switches controlled by independent OpenFlow controllers.  $c1$  and  $c2$  share and synchronize state information via an independent controller network.  $s1$  is controlled by  $c1$  and  $s2$  is controlled by  $c2$  respectively.

messages generated do not alter the existing topology. Indeed, receiving a *Packet-in* for a flow that exists in the switch is suspicious but we are not aware of any work that keeps track of such events.

3) *Switch Identification*: We now demonstrate how two malicious switches may teleport their presence using switch identification. We set up Mininet and ONOS as shown in Figure 7 with only  $s1$ ,  $s3$  and  $s4$  having connected to  $c0$  with DPID 1, 3 and 4 respectively. Also, there are no flows installed on the switches for hosts to communicate. We modified the Mininet script to configure  $s2$  with the same DPID as  $s1$ .

When  $s2$  tries to connect to  $c0$  with DPID 1 after  $s1$  has connected to  $c0$ , it is denied a connection. This way,  $s1$  teleports its presence to  $s2$ .

In Floodlight and OpenDaylight, when  $s2$  attempts to connect to  $c0$  with DPID 1 after  $s1$  has connected, Floodlight terminates the connection with  $s1$  and accepts  $s2$ 's connection.  $s2$  thereby teleports its presence to  $s1$ . Interestingly RYU allowed switches with the same DPID to co-exist which potentially introduces additional issues.

Switch identification teleportation is also possible when multiple controllers manage independent switches. We set up Mininet and ONOS as shown in Figure 8. Initially  $s1$  connects to  $c1$  with DPID 1.  $c1$  then declares itself as the *Master* for  $s1$ . At a later time,  $s2$  connects to  $c2$  and claims to have DPID 1.  $c2$  then sends  $s2$  the *Equal* role. In this manner,  $s1$  teleports its presence to  $s2$ . By inspecting the OpenFlow channels, we verified the different *Role-request* messages sent by the respective controllers to their respective switches.

#### D. Exfiltration

Our next attack is related to data exfiltration. This is a key concern for many organizations that own intellectual property, personal data or any kind of sensitive information. Once an attacker gets into a network, one possible goal of the attacker is to stealthily exfiltrate sensitive data.

We demonstrate exfiltration by considering a scenario where a small number of hosts are networked together in a remote




Fig. 9: An SDN topology with OpenFlow switches  $s1$ ,  $s2$ ,  $s3$  and  $s4$  and an OpenFlow controller  $c0$ .  $k1$  and  $k3$  are connected to  $s1$  while  $k2$  and  $k4$  are connected to  $s2$ . Note that  $s2$  is not connected to the other switches, and thereby is isolated in the data plane.  $k2$  can still exfiltrate data to  $k1$  using *out-of-band forwarding teleportation* circumventing the data plane isolation.

location. The data plane isolation is meant to improve security. However the data plane elements are managed by a controller that handles other similar remote locations. We show that in such a network, not only malicious switches can exfiltrate data using out-of-band forwarding teleportation but even malicious hosts.

We set up Mininet and ONOS as shown in Figure 9. ONOS has the *ifwd* application activated. By showing how  $k2$  can exfiltrate data to  $k1$ , we also demonstrate how  $s2$  can also exfiltrate data to  $k1$  or  $s1$ .

Given that  $s1$  and  $s2$  do not have flow rules for traffic from  $k2$  to  $k1$  (as they are located in disconnected data planes),  $k2$  can exfiltrate data to  $k1$  by simply sending a packet (e.g., UDP packet) to  $k1$  thereby exploiting out-of-band forwarding teleportation. The controller will receive the packet from  $s2$  and send it to  $s1$  which will then forward the packet to  $k1$ .

By inspecting the OpenFlow channels, we can see the out-of-band forwarding teleportation, first as a *Packet-in* and then as a *Packet-out*.

#### E. Evading Policy Conflicts

For an attacker, remaining stealthy is key to persistent existence. One of the side effects of using the out-of-band forwarding teleportation is the *Flow-mod* messages issued by the controller. The *Flow-mod* messages may generate policy conflicts (unauthorized/conflicting flow rules), alerting the administrator. A stealthier version of using the out-of-band forwarding teleportation would be to prevent the *Flow-mod* side effect. This would not only prevent policy conflicts, but also leave minimal traces on the source and sink switches.

In order to demonstrate this attack, we set up Mininet and ONOS with *ifwd* activated as shown in Figure 9.  $k2$  can exfiltrate data to  $k1$  using out-of-band forwarding teleportation without triggering *Flow-mod*'s on  $s2$  and  $s1$  by masquerading






Fig. 10: The message sequence pattern for evading policy conflicts using out-of-band forwarding teleportation. The side effect of *Flow-mod* messages are avoided when Jumbo frames are used from a masqueraded MAC address; only *Packet-ins* and *Packet-outs* are used.

its source MAC address and `ETHER_TYPE` (e.g., Jumbo frame: `0x8870`).

If the packet processor and intent framework cannot correctly identify a packet, their behavior may violate security policies. Note that it is enough if the `ETHER_TYPE` is set to a value that ONOS does not recognize, and we are not restricted to Jumbo frames only. The message sequence pattern for out-of-band forwarding teleportation without the *Flow-mod* side effect is shown in Figure 10.

By inspecting the OpenFlow channels, we can verify that the packet was indeed teleported via out-of-band forwarding teleportation first as a *Packet-in* and then as a *Packet-out*. By inspecting the flows on the switches, we can verify that no new flows are present.

*Remark on a Denial-of-service Attack:* Interestingly, we observed that a side effect of our out-of-band forwarding teleportation is a novel denial-of-service attack. If in our evading policy conflicts example, the host sends the same packet (Jumbo frame) again, then *ifwd* encounters a null-pointer exception and disconnects the switch that sent it the packet. This shows how a malicious host can cause the switch it is connected to, to be disconnected from the controller even when a packet it sends is not corrupted.

We emphasize that this is a side effect of out-of-band forwarding teleportation only, and not a teleportation issue in itself. Fortunately, the issue has been resolved by the open-source community after we contacted them (published as CVE-2015-7516).

#### F. Man-In-The-Middle

While we have so far focused on attacks where either only switches or only hosts are malicious, we now detail an attack that involves a malicious switch and a malicious host. The damage of such a collaboration can be severe, for example, the attackers could serve benign hosts with malicious web pages. In order to exemplify the attack we use HTTP rather than HTTPS.

For this attack, we set up Mininet and ONOS with *ifwd* activated as shown in Figure 11. *s1* and *k2* are both malicious while the others are not. *k3* is a benign web server. *s1* teleports specific HTTP traffic towards *k2*. *k2* modifies the HTTP traffic and teleports it back to *s1* who then forwards it to *k1*. In order to emulate the malicious switch, we introduced a flow rule




Fig. 11: An SDN topology with OpenFlow switches *s1* and *s2* with *c0* the OpenFlow controller. *k1* and *k3* are connected to *s1* while *k2* is connected to *s2*. *fw1* denies *k2* to communicate with *k1* and *k3* and vice-versa via the data plane. *s1* and *k2* being malicious, exploit the out-of-band forwarding teleportation to eavesdrop and modify communication data between *k1* and *k3* bypassing *fw1*.

(shown in Listing 1) that rewrites the destination MAC address for TCP traffic with PSH and ACK flags sent from *k3* to *k1* to *k2*. This modified packet is then passed through the flow table lookup again by using the `resubmit` action in Open vSwitch. *k2* runs *ettercap* to modify the TCP/HTTP payload and forwards the packet to the correct destination. Specifically, we created an *ettercap* filter that looks inside HTTP responses from *k3* for the word “good”, replaces it with “evil” and sends it to *k1*. The firewall *fw1* is meant to block traffic between hosts on the right and the left.

When *k1* requests the `index.html` page from *k3*, based on the flow rule installed on *s1*, only HTTP responses from *k2* are teleported to *s2* and forwarded to *k2*, through the out-of-band forwarding teleportation. Subsequently, *k2* modifies only the `index.html` web page and has *s2* teleport it back to *s1* via out-of-band forwarding teleportation. Indeed, the side effect is *Flow-mod* messages to *s1* and *s2*.

By viewing the `index.html` file received at *k1* we verified that the mitm attack was successful. The benign and malicious web pages are shown in Listing 2 and Listing 3 respectively. By inspecting the flow counters on the switches we verified that necessary packets did not pass through the data plane.

Note that we did not introduce code into the Open vSwitches to handle the mitm, therefore once the flows are installed on the switches, the firewall will block all traffic between *s1* and *s2* and vice-versa.

## VII. COUNTERMEASURES

Having showcased a variety of attacks using teleportation, we now start exploring possible countermeasures. Although we have demonstrated all the attacks using ONOS we believe that these issues are likely to become more general in nature. They are becoming important with the shift towards automated

```
priority=50001,tcp,in_port=2,
dl_src=00:00:00:00:00:03,
dl_dst=00:00:00:00:00:01,
tp_src=80,tcp_flags+=psh+ack
actions=mod_dl_dst:00:00:00:00:00:02,
resubmit:0
```

Listing 1: An Open vSwitch flow rule that was introduced into the malicious switch (*s1*) to teleport HTTP traffic with the PSH and ACK flags to the benign switch *s2*. The matching packets have the destination MAC address modified and resubmitted to the flow-table lookup which results in *packet teleportation*.

```
root@Mininet-vm:~# curl http://10.0.0.2
<html>
<head>
<title>Welcome page</title>
<body>
good
</body>
</html>
```

Listing 2: HTML code from the benign web server. Note the word “good” is present in the body of the HTML code.

and intent aware controller frameworks allowing for simpler and agnostic controller applications. Based on our experiments we have also seen that the resources required and utilized for teleportation, even at high rates are low. Therefore, it may be difficult to distinguish the attack traffic from the benign traffic. Accordingly we believe that, with the separation of the control and data plane, it is now important to monitor and police the communication channel between the separated planes.

#### A. Packet-in-Packet-out Watcher

In order to prevent the out-of-band forwarding teleportation, we strongly advise the use of a *Packet-in* and *Packet-out* watcher. It can either exist as a controller application or as an application that resides between the controller and switches akin to hypervisors. It would involve tracking and enforcing security policies for *Packet-ins* and their corresponding *Packet-outs*. Existing security enforcement kernels, hypervisors and security applications must account for *Packet-ins* and *Packet-outs* in addition to *Flow-mods* to detect and prevent out-of-band forwarding teleportation.

Note that the out-of-band forwarding teleportation in essence could also be used by malicious controller applications. In a non-adversarial scenario, the order in which a packet’s fate is decided upon by various applications can inadvertently teleport the packet. Therefore, verifying that the *Packet-out* does not reach an undesired switch/host can prevent out-of-band forwarding teleportation.

#### B. Audit-Trails and Accountability

We propose controllers to introduce secure audit-trail capabilities and accounting that enable network administrators to thoroughly investigate events in their networks. For example, controllers must log and alert sensitive events such a moving

```
root@Mininet-vm:~# curl http://10.0.0.2
<html>
<head>
<title>Welcome page</title>
<body>
evil
</body>
</html>
```

Listing 3: HTML code modified by the malicious switch *s1* and host *k2*. Note the word “evil” is present in the body of the HTML code.

MAC addresses or receiving a *Packet-in* when a flow has not yet timed out. Such capabilities can aid detection and prevention mechanisms. It is also useful for investigating security incidents. We recommend administrators to frequently view controller logs, investigate failed events and suspicious identities in the network.

#### C. Enhanced IDS with Waypoint Enforcement

Network intrusion detection systems are an important means to detect and limit cyber attacks today, and accordingly intrusion detection systems constitute an integral part of most networks. We strongly suggest the use of an IDS application on top of the controller, that can inspect *Packet-ins* and *Packet-outs* and alert on suspicious traffic. Indeed, some controllers today already offer basic functionality for waypoint enforcement. In particular, we suggest waypoint enforcement and coordinating intrusion detection systems from the control plane with the data plane. This is non-trivial, but vital for network security.

### VIII. RELATED WORK

While researchers have already pointed out several interesting novel challenges in providing a correct operation of networks with separate data and control planes [31], [34], [39], it is generally believed that SDN has the potential to render computer networking more verifiable [20], [21] and even secure [18], [28], [29], [33], [36], [37].

Only recently researchers have started discovering security threats in SDN. Kloti et al. [22] report on a STRIDE threat analysis of OpenFlow, and demonstrate data plane resource consumption attacks. Kreutz et al. [23] survey several threat vectors that may enable the exploitation of SDN vulnerabilities. A key challenge arising from the separation of the control and data planes is the potential loss of network visibility. It has been shown that the network view of the controller may even be poisoned [11], [15].

While much research went into designing more robust and secure control planes [9], [10], less published work exists on the issue of malicious switches. A notable exception is the work by Antikainen et al. [7], who consider the possibility of a malicious relay node for a man-in-the-middle attack. Interestingly, in our paper, we have shown that the relay node can be the benign controller itself.

To the best of our knowledge, our work is the first to point out and characterize the fundamental problem of SDN

teleportation. More generally, while most prior studies about malicious switches focus on (indirect) *attacks targeting the controller*, we in this paper demonstrate new kinds of attacks which merely exploit the controller for *directly* attacking (e.g., the availability or confidentiality) of network services.

However, there are a number of interesting approaches proposed in the literature which have implications for our scenarios as well. For example, the pre- and post-conditions of Topoguard [15] can defend against our path update attack. However, if the switches are malicious, these conditions can be spoofed by the malicious switches. Also, Topoguard cannot detect teleportation using path reset, switch identification and out-of-band forwarding teleportation.

Sphinx [11] can alert on the path update teleportation. However, it cannot detect the path reset as the flow graph remains the same. Additionally, Sphinx assumes that switches cannot use the same DPIDs therefore, we believe that our switch identification teleportation will not be detected by Sphinx. Also, our out-of-band forwarding relies on *Packet-in* and *Packet-out* messages, while *Packet-outs* are not considered by Sphinx<sup>2</sup>. Therefore the suggested out-of-band forwarding teleportation can evade Sphinx, until topology altering flows are installed.

Porras et al. [32] propose a security mediator that comprises of Rule Conflict Analysis, Role-based Source Authentication, State Table Manager and a Permission Mediator. We admit that the path update can be detected using this approach, however, our path reset does not introduce any conflicting rules. The *Features-reply* messages are not a part of their solution therefore, we believe that *switch identification* teleportation can succeed. With respect to out-of-band forwarding teleportation, unless the mediator investigates the destination switch or MAC address in the *Packet-out*, the teleportation can bypass the security mediator given sufficient permissions.

SDN Hypervisors such as CoVisor [16], Flowvisor [35], FortNOX [33] depend on policies maintained in the hypervisor. Therefore, we believe that all our teleportation mechanisms hold unless a specific policy blocks it.

Dover Networks [12] discovered the behavior of Floodlight with switches using the same DPID, which we exploit for teleportation.

While *Security-Mode ONOS* [17] can enhance the security in many scenarios, by introducing roles and permissions, at least today, it does not help against teleportation: once *ifwd* has the permission to write intents and emit packets, our teleportation succeeds. These permissions are bare necessities for *ifwd* to function.

## IX. CONCLUSIONS

As OpenFlow networks transition from research to production, new levels of reliability and performance are necessary [24]. This paper has identified and demonstrated a novel security threat introduced by software-defined networks separating the control plane from the data plane, in the presence of an unreliable south-bound interface (containing malicious

Attack	Teleportation technique	Exploited by
Bypass Firewall	Out-of-band forwarding	Switch and Host
Bypass NIDS	Out-of-band forwarding	Switch and Host
Rendezvous	Path update Path reset Switch identification	Switch Switch Switch
Exfiltration	Out-of-band forwarding	Switch and Host
Evading policy conflicts	Out-of-band forwarding	Switch and Host
Man-in-the-middle	Out-of-band forwarding	Switch and Host

TABLE II: Summary of teleportation attacks and involved entities.

switches): we have shown that state-of-the-art controller(s) are vulnerable to teleportation. Teleportation has numerous applications (cf. the summary in Table II): it can be exploited to bypass security-critical network elements (e.g., to exfiltrate confidential information), as a discovery protocol for malicious switches, to evade policy conflicts as well as for man-in-the-middle attacks. Based on our preliminary evaluation, we can say that even a teleportation channel of over 10 Mbps can easily be used inside a loaded control channel.

Our work can also be seen as a first security analysis of the increasingly popular intent-based network mechanisms: while intent-based mechanisms are attractive for allowing (cloud) network operators resp. SDN applications to focus on “what to connect” rather than “how”, we have shown that controller managed intents need to be used with care. Indeed, our experiments with controllers that are only starting to introduce an intent based mechanism are not yet vulnerable to all the specific attacks presented in this paper. Moreover, while intent mechanism implementations can vary across controllers, we believe that the underlying issues are fundamental.

We understand our work as a first step, and believe that our paper opens several relevant directions for future research. In particular, we plan to extend our vulnerability analysis to other SDN protocols and conduct a more in-depth performance analysis. Another relevant avenue for future research regards the development of countermeasures.

**Acknowledgments.** Research (partially) supported by the Danish Villum project *ReNet*. We thank Jean-Pierre Seifert, Tobias Fiebig, and Jens Sieberg, for several discussions and feedback on our work.

## REFERENCES

- [1] The tale of one thousand and one dsl modems. <https://securelist.com/analysis/publications/57776/the-tale-of-one-thousand-and-one-dsl-modems/>, 2012.
- [2] Huawei hg8245 backdoor and remote access. <http://websec.ca/advisories/view/Huawei-web-backdoor-and-remote-access>, 2013.
- [3] Netis routers leave wide open backdoor. <http://blog.trendmicro.com/trendlabs-security-intelligence/netis-routers-leave-wide-open-backdoor/>, 2014.
- [4] Snowden: The NSA planted backdoors in cisco products. <http://www.infoworld.com/article/2608141/internet-privacy/snowden--the-nsa-planted--backdoors-in-cisco-products.html>, 2014.
- [5] Synful knock - a cisco router implant - part i. [https://www.fireeye.com/blog/threat-research/2015/09/synful\\_knock\\_-\\_acis.html](https://www.fireeye.com/blog/threat-research/2015/09/synful_knock_-_acis.html), 2015.
- [6] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the internet impasse through virtualization. *Computer*, 38(4):34–41, Apr. 2005.

<sup>2</sup> Unfortunately, the source code of Sphinx is not available.

- [7] M. Antikainen, T. Aura, and M. Sarela. Spook in your network: Attacking an sdn with a compromised openflow switch. In *Secure IT Systems*, volume 8788 of *Lecture Notes in Computer Science*. 2014.
- [8] austinmarton. Send a raw ethernet frame in linux. <https://gist.github.com/austinmarton/1922600>, 2011.
- [9] M. Canini, P. Kuznetsov, D. Levin, and S. Schmid. A distributed and robust sdn control plane for transactional network updates. In *Proc. IEEE INFOCOM*, 2015.
- [10] M. Canini, D. Venzano, P. Perešini, D. Kostić, and J. Rexford. A nice way to test openflow applications. In *Proc. 9th USENIX NSDI*, 2012.
- [11] M. Dhawan, R. Poddar, K. Mahajan, and V. Mann. Sphinx: Detecting security attacks in software-defined networks. In *Proc. NDSS*, 2015.
- [12] J. M. Dover. A denial of service attack against the open floodlight sdn controller. In *Technical Report*, 2013.
- [13] S. Fagerland, W. Grange, and Blue Coat Systems Inc. The inception framework: Cloud-hosted APT. [http://dc.bluecoat.com/Inception\\_Framework](http://dc.bluecoat.com/Inception_Framework), 2014.
- [14] I. S. W. Group. Internet as a critical infrastructure: Security, resilience and dependability aspects (jra7). <http://www.internet-science.eu/groups/internet-critical-infrastructure-security-resilience-and-dependability-aspects>, 2015.
- [15] S. Hong, L. Xu, H. Wang, and G. Gu. Poisoning network visibility in software-defined networks: New attacks and countermeasures. In *Proc. NDSS*, 2015.
- [16] X. Jin, J. Gossels, J. Rexford, and D. Walker. Covisor: A compositional hypervisor for software-defined networks. In *Proc. USENIX NSDI*, 2015.
- [17] P. Joshi. ONOS wiki home. <https://wiki.onosproject.org/display/ONOS/Security-Mode+ONOS>, 2015.
- [18] P. Kampanakis, H. Perros, and T. Beyene. Sdn-based solutions for moving target defense network protection. In *Proc. IEEE WoWMoM*, 2014.
- [19] N. Kang, Z. Liu, J. Rexford, and D. Walker. Optimizing the "one big switch" abstraction in software-defined networks. In *Proc. ACM CoNEXT*, 2013.
- [20] P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. In *Proc. 9th USENIX NSDI*, 2012.
- [21] A. Khurshid, X. Zou, W. Zhou, M. Caesar, and P. B. Godfrey. Veriflow: Verifying network-wide invariants in real time. In *Proc. 10th USENIX NSDI*, 2013.
- [22] R. Kloti, V. Kotronis, and P. Smith. Openflow: A security analysis. In *Proc. IEEE ICNP*, Oct 2013.
- [23] D. Kreutz, F. M. Ramos, and P. Verissimo. Towards secure and dependable software-defined networks. In *Proc. ACM HotSDN*, 2013.
- [24] M. Kuniar, P. Pereni, and D. Kostić. What you need to know about sdn flow tables. In *Proc. Passive and Active Measurement (PAM)*. 2015.
- [25] S. Lee, T. Wong, and H. S. Kim. Secure split assignment trajectory sampling: A malicious router detection system. In *International Conference on Dependable Systems and Networks (DSN'06)*, pages 333–342. IEEE, 2006.
- [26] F. Lindner. Cisco ios router exploitation. *Black Hat USA*, 2009.
- [27] M. Armbrust et al. A view of cloud computing. *Commun. ACM*, 53(4):50–58, Apr. 2010.
- [28] S. A. Mehdi, J. Khalid, and S. A. Khayam. Revisiting traffic anomaly detection using software defined networking. In *Proc. RAID*, pages 161–180, 2011.
- [29] J. Mudigonda, P. Yalagandula, J. Mogul, B. Stiekes, and Y. Pouffary. Netlord: A scalable multi-tenant network architecture for virtualized datacenters. In *Proc. ACM SIGCOMM*, pages 62–73, 2011.
- [30] P. v. Oorschot, T. Wan, and E. Kranakis. On interdomain routing security and pretty secure bgp (psbgp). *ACM Trans. Inf. Syst. Secur.*, 10(3), July 2007.
- [31] O. Padon, N. Immerman, A. Karbyshev, O. Lahav, M. Sagiv, and S. Shoham. Decentralizing sdn policies. In *Proc. ACM POPL*, 2015.
- [32] P. Porras, S. Cheung, M. Fong, K. Skinner, and V. Yegneswaran. Securing the software-defined network control layer. In *Proc. NDSS*, 2015.
- [33] P. Porras, S. Shin, V. Yegneswaran, M. Fong, M. Tyson, and G. Gu. A security enforcement kernel for openflow networks. In *Proc. ACM HotSDN*, pages 121–126, 2012.
- [34] M. Reitblatt, N. Foster, J. Rexford, and D. Walker. Consistent updates for software-defined networks: Change you can believe in! In *Proc. HotNets*, pages 7:1–7:6, 2011.
- [35] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar. Flowvisor: A network virtualization layer. *OpenFlow Switch Consortium, Tech. Rep.*, 2009.
- [36] S. Shin, P. Porras, V. Yegneswaran, M. Fong, G. Gu, and M. Tyson. Fresco: Modular composable security services for software-defined networks. In *Proc. NDSS*, 2013.
- [37] S. Shin, V. Yegneswaran, P. Porras, and G. Gu. Avant-guard: Scalable and vigilant switch flow management in software-defined networks. In *Proc. ACM CCS*, pages 413–424, 2013.
- [38] System Administrator. ONOS wiki home. <https://wiki.onosproject.org/display/ONOS/ONOS+Wiki+Home>, 2015.
- [39] W. Zhou, D. K. Jin, J. Croft, M. Caesar, and P. B. Godfrey. Enforcing customizable consistency properties in software-defined networks. In *Proc. USENIX NSDI*, pages 73–85, 2015.